

TESLA: An Energy-saving Agent that Leverages Schedule Flexibility

Jun-young Kwak, Pradeep Varakantham*, Rajiv Maheswaran, Yu-Han Chang,
Milind Tambe, Burcin Becerik-Gerber, Wendy Wood
University of Southern California, Los Angeles, CA, 90089
*Singapore Management University, Singapore, 178902
{junyounk,maheswar,ychang,tambe,becerik,wendy.wood}@usc.edu,
*pradeepv@smu.edu.sg

ABSTRACT

This innovative application paper presents TESLA, an agent-based application for optimizing the energy use in commercial buildings. TESLA's key insight is that adding flexibility to event/meeting schedules can lead to significant energy savings. TESLA provides three key contributions: (i) three online scheduling algorithms that consider flexibility of people's preferences for energy-efficient scheduling of incrementally/dynamically arriving meetings and events; (ii) an algorithm to effectively identify key meetings that lead to significant energy savings by adjusting their flexibility; and (iii) surveys of real users that indicate that TESLA's assumptions exist in practice. TESLA was evaluated on data of over 110,000 meetings held at nine campus buildings during eight months in 2011–2012 at USC and SMU. These results show that, compared to the current systems, TESLA can substantially reduce overall energy consumption.

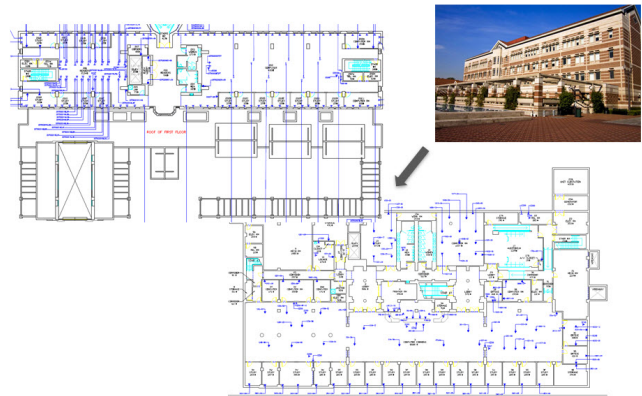


Figure 1: The actual research testbed (library) at USC

Categories and Subject Descriptors

I.2.11 [ARTIFICIAL INTELLIGENCE]: Distributed Artificial Intelligence

General Terms

Algorithms, Experimentation, Human Factors

Keywords

Innovative Applications, Energy, Sustainable Multiagent Building Application, Energy-oriented Scheduling

1. INTRODUCTION

Reducing energy consumption is an important goal for sustainability. Thus, conserving energy in commercial buildings is important as it is responsible for significant energy consumption. In 2008, commercial buildings in the U.S. consumed 18.5 QBTU, representing 46.2% of building energy consumption and 18.4% of U.S. energy consumption [10]. This energy consumption is significantly affected by a large number of meetings or events in those buildings. Furthermore, a recent study shows that meeting frequency in commercial buildings is significant and continues to grow [4]. In 2001,

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA. Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

U.S. Fortune 500 companies are estimated to have held 11 million formal meetings daily and 3 billion meetings yearly.

Energy-oriented scheduling can assist in reducing such energy consumption [5, 15, 19]. Although conventional scheduling techniques compute the optimal schedule for many meetings or events while satisfying their given requirements (i.e., computing a valid schedule) [8, 13], they have not typically considered energy consumption explicitly. More recently, there have been some trials to conserve energy by consolidating meetings in fewer buildings [2, 3]. In particular, Portland State University consolidated night and weekend classes, which are previously scattered across 21 buildings, into five energy efficient buildings. By doing this, they reported that electricity consumption was reduced by 18.5% (78,000 kWh) in the autumn compared to the previous three-year average. Similarly, Michigan State University consolidated classes and events into fewer buildings on campus, and energy reductions in the seven buildings ranged from 2–20%, saving \$16,904. However, these efforts have been decided manually, and no underlying intelligent system was used.

Motivated by this prior work, we describe TESLA (Transformative Energy-saving Schedule-Leveraging Agent), an innovative agent-based application for optimizing the use of facilities in commercial buildings. TESLA's key insight is that adding flexibility to meeting schedules can lead to significant energy savings. TESLA provides three key contributions. First, it provides three scheduling algorithms — myopic, predictive non-myopic, and full-knowledge optimization — that consider *flexibility* of people's preferences for energy-efficient scheduling of incrementally/dynamically arriving meetings and events. In this work, flexibility specifically refers to



Figure 2: The current room reservation system at the testbed building

the number of options made available by the scheduling constraints in terms of starting time, locations and the deadline before committing to the finalized schedule details. Second, using the predictive non-myopic method, TESLA presents an algorithm to effectively identify key meetings that could lead to significant energy savings by adjusting their flexibility. Third, surveys of real users are provided indicating that TESLA's savings can be realized in practice by effectively leading people to change their schedule flexibility. To validate our work, we have used a public domain simulation testbed [10] with details in our testbed building and validated this simulation. Just within this testbed building, our results show that, in a validated simulation, TESLA is projected to save about 250 kWh of energy (roughly \$17K) annually. If this pilot is successful, TESLA can offer energy saving benefits to all commercial buildings where meetings affect energy usage.


The rest of the paper is organized as follows: In Section 2, we describe our testbed buildings along with real data from those buildings. In Section 3, we describe the TESLA system and scheduling algorithms at the heart of it. Section 4 provides evaluations of each of our algorithms using real-world meeting and energy data indicating that TESLA could potentially provide significant savings in overall energy consumption.

2. RESEARCH TESTBED

2.1 Educational Building Testbed

Our system is to be deployed in an educational building. Figure 1 shows the testbed building for TESLA's deployment and the floor plans of 2nd and basement floors. It is one of main libraries at the University of Southern California and has been designed with a building management system. It hosts a large number of meetings (about 300 unique meetings per regular day) across 35 group study rooms. Each study room has different physical properties including different types and numbers of devices and facilities (e.g., video conferencing equipment, computer, projector, video recorder, office electronic devices, etc.), room size, lighting specification, and maximum capacity (4 – 15 people). This building operates these study rooms 24 hours a day and 7 days a week except on national holidays. The temperature in group study rooms is regulated by the facility managers according to two set ranges for occupied and unoccupied periods of the day. HVAC (Heating, Ventilating, and Air Conditioning) systems always attempt to reach the pre-set temperature regardless of the presence of people and their preferences in terms of temperature. Lighting and appliance devices are manually controlled by users.

In this building, meetings are requested by users by a centralized online room reservation system (see Figure 2). In the current reservation system, no underlying intelligent system is used; instead, users reactively make a request based on the availability of room and time when they access the system. While users make a request



(a) Meeting frequency data

(b) Distribution of total meeting requests per day

Figure 3: Real data analysis

Table 1: Meeting request arrival distribution

Time period	Likelihood (%)
1 day before	55.73
1-2 days before	18.40
2-3 days	8.72
3-4 days	5.52
4-5 days	3.68
5-6 days	3.05
6-7 days	3.35
> 7 days	1.56

using the system, they are asked about additional information including the number of meeting attendees and special requirements. Reservations can be made up to 7 days in advance.

Given the significant number of meetings per day and the centralized online meeting reservation system, it provides a good environment to test various energy-oriented scheduling techniques to mitigate energy consumption. TESLA's goal is to enable users to input flexibility in their scheduling request, to identify key scheduling requests, and use this information in algorithms that can provide energy-efficient schedules to effectively conserve energy in commercial buildings. To evaluate TESLA, we have built upon a simulation testbed using real building data [10] and validated with real-world energy data.¹ This validated simulation environment is used to evaluate TESLA with real meeting data. In addition, we also test TESLA on buildings at the Singapore Management University as described below.

2.2 Data Analysis

In collaboration with building system managers, we have been collecting data specifying the past usage of group study rooms, which are collected for 8 months (January through August in 2012) at USC. The data for each meeting request includes the time of request, starting time, time duration, specified room, and group size. The data set contains 32,065 unique meetings, and their average meeting time duration is 1.78 hours.

Figure 3(a) shows the actual meeting frequency (y-axis) over time (24 hours, x-axis) of sampled 4 locations (out of 35 rooms). Figure 3(b) shows the probability distribution over total meeting requests per day. The x-axis of the figure indicates the total number of meeting requests per day (ranging from 0 to about 350) and the y-axis shows how likely the system will have the given number of total meeting requests (x-axis) on one day. One can see that the probability of having 50 or fewer meetings is 42.92% and the probability of having 250 or more meetings is 30.04%. These are used

¹Energy validation results can be found here: <http://teamcore.usc.edu/junyounk/TESLA-sp.pdf>

to estimate the model of future meetings in our algorithm that will be presented in Section 3.2.

Table 1 shows how early meeting requests were made. In the table, column 2 indicates the percentage of meetings that were requested within the given time period (column 1). For instance, 44.27% of all meeting requests were made 1 day before (or even earlier) the actual meeting day. This analysis would be helpful in understanding how our algorithm could achieve significant energy savings in this domain.

While evaluating TESLA, we also consider another data set containing over 80,000 meetings that have been collected for three months in 2011 from over 500 conference/meeting rooms across eight buildings at SMU, which gives us a sense regarding how TESLA will handle energy-oriented scheduling problems in large buildings.

3. TESLA

3.1 Application Architecture

TESLA is a goal-seeking (to save energy), continuously running autonomous agent. TESLA performs on-line energy-efficient scheduling while considering dynamically arriving inputs from users; these dynamic inputs makes the scheduling complex and TESLA needs to learn a predictive model for users' inputs and preferences. More specifically, TESLA:

- takes inputs (i.e., preferred time, location, the number of meeting attendees, etc.) from different users and their proxy agents at different times (Sections 2 & 3.1)
- autonomously performs on-line energy-efficient scheduling as requests arrive while balancing user comfort (Section 3.2.1)
- autonomously, on own initiative, interacts with different users based on identified problematic key meetings (Section 3.2.2) to persuade them to change meeting flexibility for minimizing bother cost to users
- bases its non-myopic optimization on learned patterns of meetings (Sections 3.2.1 & 4)

TESLA communicates with other users or their proxy agents, who have the corresponding meeting attendees' preference and behavior models with a certain level of adjustable autonomy [14]. Proxy agents communicate on behalf of meeting attendees with TESLA. Meeting requests are the information we get from the interface of TESLA. TESLA may also communicate with proxy agents to adjust the given meeting flexibility of key meetings to compute energy-efficient schedules. TESLA focuses on minimizing unnecessary interactions by detecting a small number of key meetings while negotiating with people to adjust their flexibility.

3.2 Algorithm

The objective of this work is to come up with energy efficient schedules in commercial buildings with a large number of meetings while considering (i) flexibility in meeting requests over time, location and deadline; and (ii) user preferences with respect to energy and satisfaction. To account for these two constraints, we provide two types of algorithms, which are at the heart of TESLA. First, we provide algorithms that compute a schedule for known and predicted meeting requests which have flexibility in time, location and deadline. Second, based on the schedule obtained, we provide algorithms that detect meeting requests which if modified (to increase flexibility) can result in significant energy savings.

3.2.1 Scheduling algorithms

Before describing our scheduling algorithms, we formally describe the scheduling problem. Let T represent the entire set of time slots available and L represent the set of available locations each day. A schedule request r_i is represented as the tuple: $r_i = \langle a_i, T_i, L_i, \delta_i, d_i, n_i \rangle$, where: a_i is the arrival time of the request, $T_i \subset T$ is the set of preferred time slots for the start of the event and $L_i \subset L$ is a set of preferred locations. d_i is the deadline by which the time and location for the meeting should be notified to the user, δ_i is the duration for the event and finally, n_i is the number of attendees.

The flexibility of the meeting request r_i , α_i , is a vector of three values: $\langle \alpha_i^T, \alpha_i^L, \alpha_i^d \rangle$.

- α_i^T : time flexibility of meeting i (%). $\alpha_i^T = \frac{|T_i|-1}{|T|-\delta_i} \times 100$ ($|T| > \delta_i$; i.e., $|T|$ is 24 hours per day). Given only one time slot ($|T_i| = 1$), $\alpha_i^T = 0$ and all available time slots ($|T_i| = |T| - \delta_i + 1$), $\alpha_i^T = 100$. For example, assuming that people give $T_i = 4-7$ pm on Monday and its meeting time duration is 2 hours, then $\alpha_i^T = (4-1)/(24-2) \times 100 = 13.64\%$.
- α_i^L : location flexibility of meeting i (%). $\alpha_i^L = \frac{|L_i|-1}{|L|-1} \times 100$ ($|L| > 1$). Given only one location slot ($|L_i| = 1$), $\alpha_i^L = 0$ and given all available locations ($|L_i| = |L|$), $\alpha_i^L = 100$.
- α_i^d : deadline flexibility of meeting i (%). $\alpha_i^d = \frac{d_i - a_i}{d_i^* - a_i} \times 100$, where d_i^* is the latest notification time (e.g., midnight on the meeting day) ($d_i^* > a_i$). $0 \leq \alpha_i^d \leq 100$

Given a set of requests, R , we provide a mixed integer linear program (MILP) to compute a schedule that minimizes the overall energy consumption (and will be used in our algorithms below). Here is the notation that will be employed in the MILP:

- $x_{i,l,t}$ is a binary variable that is set to 1 if meeting request r_i is scheduled in location l starting at time t .
- $E_{i,l,t}$ is a constant that is computed for a meeting request r_i if it is scheduled in location l at time t using the HVAC energy consumption equations.
- C is a constant that indicates the reduction in energy consumption because of scheduling a meeting in the previous time slot.
- $e_{i,l,t}$ is a continuous variable that corresponds to the energy consumed because of scheduling meeting i in location l at time t . The value of this variable is affected based on whether there is a meeting scheduled in the previous time slot ($t - 1$). $e_{i,l,t} = E_{i,l,t} - x_{i,l,t-1} \cdot C$ ²
- $S_{i,l,t}$ is an value that indicates the satisfaction level obtained with users in meeting request r_i for scheduling the meeting in location l at time t . B is a threshold on the satisfaction level required by users.
- M is an arbitrarily large positive constant.

² $e_{i,l,t}$ gets affected by a meeting in the previous time slot in the same location. This is because adjacent meetings affect the indoor temperature, which makes HVACs operate differently to maintain the desired temperature level.

$$\min \sum_{t \in T} \sum_{l \in L} \sum_{i \in R} x_{i,l,t} \cdot e_{i,l,t} \quad (1)$$

s.t.

$$e_{i,l,t} = E_{i,l,t} - x_{i,l,t-1} \cdot C, \quad \forall i \in R, l \in L, t \in T \quad (2)$$

$$\sum_{t \in T} \sum_{l \in L} x_{i,l,t} \cdot S_{i,l,t} \geq B, \quad \forall i \in R \quad (3)$$

$$\sum_{i \in R} x_{i,l,t} \leq 1, \quad \forall l \in L, t \in T \quad (4)$$

$$\sum_{i' \in R \setminus \{i\}} \sum_{t' = t}^{t + \delta_i - 1} x_{i',l,t'} \leq M(1 - x_{i,l,t}), \forall l \in L, i \in R, t \in T \quad (5)$$

$$x_{i,l,t} \in \{0, 1\}, \quad \forall i \in R, l \in L, t \in T \quad (6)$$

The objective of the MILP above is finding the optimal assignment of meeting requests to locations and time slots that is characterized by the solution, $x_{i,l,t}^*$ in order to minimize energy consumption. The constraint (2) is for computing energy consumption considering the back-to-back meeting effect. The constraint (3) is for checking if the computed schedule maintains the given comfort level B . The constraints (4) and (5) are the allocation restrictions that for each assignment slot, only one meeting can be scheduled considering the given time duration of meeting.

We now define specific disjoint sets of meeting requests, R that will enable us to characterize two type of scheduling algorithms, where t is time to schedule a given set of requests, R .

- $R^S(t) = \{i : d_i = t \text{ and } a_i \leq t\}$: a set of requests that have to be scheduled at time t
- $R^A(t) = \{i : d_i < t \text{ and } a_i < t\}$: a set of requests that were assigned before time t
- $R^K(t) = \{i : d_i > t \text{ and } a_i \leq t\}$: a set of known future requests, which arrived before time t , but will be scheduled in the future
- $R^U(t) = \{i : d_i > t \text{ and } a_i > t\}$: a set of unknown future requests




Figure 4: Disjoint sets of R

As a simple example (shown in Figure 4), let us consider that we have 4 meeting requests (r_1, r_2, r_3 , and r_4), which are supposed to be scheduled on the same day. The current time is t . According to the definition, $R^S(t) = \{r_2\}$, $R^A(t) = \{r_1\}$, $R^K(t) = \{r_3\}$, and $R^U(t) = \{r_4\}$.

The MILP is the core of the three scheduling algorithms, and the algorithms call the MILP formulation using the function

GETMINENERGYSCHEDULE(R) on the actual input R and provide different results as described below.

Myopic optimization algorithm: We have the myopic optimization algorithm, which obtains a schedule by making the following function call: GETMINENERGYSCHEDULE($R^A(t) \cup R^S(t) \cup R^K(t)$). A schedule and energy consumption are obtained without accounting for future unknown meetings.

Predictive non-myopic optimization: We have a predictive optimization method that minimizes expected reduction of energy over possible unknown meetings. Let U be a set that contains various possibilities of unknown meeting requests (obtained by analyzing previous meeting data such as the one in Figure 3(b)) that could arise in future times. More specifically, we have a probability distribution over the possible number of total meeting requests per day (shown in Figure 3(b)). Then, the likelihood that k more meetings will arrive on the same day, $p^U(k)$, is computed considering that we currently have s meetings so far. For those k future meeting requests in U , we generate random request tuples (specifically, T_i & L_i) based on the distribution over the assignment spots as shown in Figure 3(a). A predictive optimization approach thus solves the following optimization problem:

$$\min_x \sum_{U \in U} p^U \cdot \text{GetMinEnergySchedule}(R^A(t) \cup R^S(t) \cup R^K(t) \cup U)$$

Full-knowledge Optimization: As a benchmark algorithm for comparison purposes, we provide the full-knowledge optimization method. In this method, assuming that the entire set of meeting requests R is given, which is ideal, we compute the final schedule using the MILP used in the myopic optimization technique. The performance comparison results will be provided in Section 4.1.

3.2.2 Identifying key meetings

The scheduling agent computes the optimal schedule considering the given flexibility of meetings. It can obtain more energy-efficient schedules by relaxing those constraints.

We provide an algorithm that finds key meetings that can reduce significant energy consumption if made more flexible.

Algorithm 1 IDENTIFYKEYMEETINGS

```

1:  $U \leftarrow \emptyset$ 
2: {Initialize a set of key meetings}
3:
4: for  $i = 1 \dots N \in R$  do
5:   { $R$  is a set of requests.}
6:    $V_i \leftarrow \text{CALEXPENERGYSAVINGS}(\alpha_i, \{\alpha'_{i,1}, \dots, \alpha'_{i,k}\})$ 
7:   { $\alpha_i$  is an initially given flexibility of meeting  $i$ , and  $\alpha'_{i,k}$  is one of
   the desired flexibility options for meeting  $i$ }
8:
9:   if  $V_i > \tau_i$  then
10:    { $\tau_i$  is a threshold for meeting request  $i$ .}
11:     $U \leftarrow U \cup \{i\}$ 
12:
13: return  $U$ 

```

We first initialize a set that will contain key meetings identified by our algorithm (line 1). For each meeting request r_i , we then compute the expected energy savings of the meeting i when its flexibility level is changed from the initial level α_i to the desired level α'_i assuming the other meetings' flexibility levels are fixed (line 6). The expected energy saving value of meeting i , $V_i = (E_{\alpha_i} - E_{\alpha'_i})/E_{\alpha_i}$ ($0 \leq V_i \leq 1$), where E_{α_i} is the current total energy consumption with the given level of flexibility α_i , and $E_{\alpha'_i}$ is the reduced total energy consumption if the meeting i 's




Figure 5: Energy savings

flexibility is changed to one of k possible options, $\alpha'_{i,k}$, while others keep their given flexibility levels. In this work, we consider a heuristic to set the threshold value to tell whether or not the current meeting i is a key meeting: a fixed single threshold value (e.g., 0.4 as a universal threshold).

4. EMPIRICAL VALIDATION

We evaluate the performance of TESLA and experimentally show that it can conserve energy by providing more energy-efficient schedules in commercial buildings. At the end of this section, we provide actual survey results that we have conducted on schedule flexibilities of real users. The experiments were run on Intel Core2 Duo 2.53GHz CPU with 8GB main memory. All techniques were evaluated for 100 independent trials and we report the average values.

4.1 Simulation Results

4.1.1 Does flexibility help?

As an important first step in deploying TESLA, we first verified if the agent could save more energy with more flexibility while scheduling given meeting and event requests. To that end, we compared energy consumption of three different approaches using the real-world meeting data mentioned in Section 2.2: (i) the current benchmark approach in use at the testbed building; (ii) a random method that randomly assigns time and location for meetings; and (iii) the optimal method using the full-knowledge optimization technique described in Section 3.2.

Figure 5 shows the average daily energy consumption in kWh computed based on schedules from the three algorithms above. In the figure, the consumption is the amount of energy consumed based on the past schedules obtained from the current manual reservation system, which shows a very similar performance to the random approach. The optimal method assuming the full amount of flexibility (i.e., 24 hours for α^T , 35 rooms for α^L and delay the deadline before which the final schedule should be informed for α^d) achieved statistically significant energy savings of 48.08% compared to the current energy consumption at the testbed site (t-test; $p < 0.01$), which is equivalent to annual savings of about \$17,600 considering an energy rate of \$0.193/kWh [1] and CO_2 emissions from the energy use of 5.5 homes for one year.

4.1.2 Online scheduling methods with flexibility:

We then compared solution qualities of the three scheduling al-

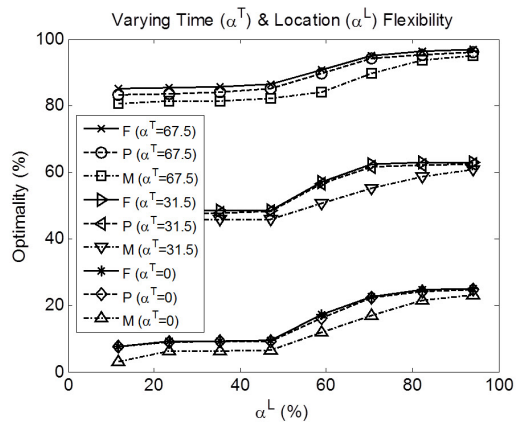


Figure 6: Energy savings while varying flexibility

gorithms in TESLA presented in Section 3.2.1. Figure 6 and Table 2 show that how much each algorithm saves when compared to the optimal value (i.e., full-knowledge optimization assuming the full flexibility) while varying the time and location flexibility level (assuming 0% deadline flexibility). *The flexibility in our model represents a 3-dimensional space (time, location and deadline), which we have fully explored. Due to the limitation of space, however, we have only shown one slice of the table with the deadline flexibility of 0%.* We show results exploring deadline flexibility later.

The optimality is computed as follows: $(E_a - E_c)/(E_a - E_o)$, where E_a is the actual energy consumption without any flexibility, E_o is the optimal energy consumption, and E_c is the computed energy consumption using three different scheduling algorithms that we compare using the real meeting data.

Figure 6 shows the average optimality in percentage of each algorithm (M: myopic, P: predictive non-myopic and F: full-knowledge) while varying the location flexibility (α^L ; x-axis) and time flexibility (α^T ; each graph assumed the different amount of α^T as indicated in the legend). In the figure, for each pair of flexibility values (α^T, α^L), we report the average optimality in percentage (i.e., 100% indicates the optimal value, and 0% means that there was no improvement from the actual energy consumption). Thus, higher values indicate better performance.

As shown in the figure, as more flexibility is given to the system, the agent can compute schedules with less energy consumption. The gain in optimality from myopic to predictive non-myopic is because the latter can leverage user flexibility to put a meeting in a suboptimal spot at the meeting request time to account for future meetings, yielding better results at the actual day of meetings. For example, a flexible meeting request can be moved away from a known popular time-location spot. We conclude that (i) the predictive non-myopic method is superior to the myopic method; (ii) the predictive non-myopic method performs almost as well as the full-knowledge optimization (about 98%); and (iii) the full flexibility is not required to start accruing benefits of flexibility.

In the real-world, it is hard to imagine that all people will simply comply to change their flexibility to achieve such optimality. Thus, we provide one additional result shown in Table 2 which varies the percentage of meetings that will have flexibility (p_f). In particular, the value of row 10 and column 5 (highlighted in the table shows the optimality achieved by the predictive method assuming that 20% of meetings (randomly selected) have $(\alpha^T, \alpha^L) = (0\%, 23.5\%)$ flexibility and the remaining 80% have no flexibility. Our main conclusions are: (i) if we increase p_f , we are able to achieve a higher optimality; and (ii) flexibility in a small number of meetings

Table 2: % of optimal energy savings: varying α^T , α^L , and p_f

		Alg.		Location flexibility (α^L)				
				23.5	47.1	70.6	94.1	
T. flex. (α^T)	0	M	1.0	6.2	6.5	17.0	23.0	
			0.8	5.3	5.6	14.5	20.5	
			0.5	4.4	4.6	12.9	18.1	
			0.2	2.8	2.9	7.7	11.3	
		P	1.0	8.9	9.3	22.2	24.7	
			0.8	7.8	8.1	20.3	22.5	
			0.5	5.8	6.0	15.4	18.1	
			0.2	3.8	3.9	9.3	12.1	
		F	1.0	9.1	9.5	22.6	25.1	
			0.8	7.8	8.5	20.2	23.1	
			0.5	6.3	6.6	16.4	18.2	
			0.2	4.0	4.5	11.1	13.3	
	31.5	M	1.0	45.4	45.9	55.3	60.7	
			P	1.0	47.5	48.2	61.5	62.5
		F	1.0	48.4	48.6	62.3	63.1	
			0.8	41.3	42.6	55.1	56.9	
			0.5	29.7	30.5	43.5	44.0	
			0.2	15.8	16.6	25.9	26.2	
		67.5	M	1.0	81.3	82.2	89.6	95.1
				P	1.0	83.6	85.1	94.2
			F	1.0	85.5	86.4	95.1	96.9
				0.8	73.1	73.0	87.1	90.4
				0.5	53.5	54.3	64.9	67.0
				0.2	29.2	30.3	37.9	40.8

(M: myopic, P: predictive non-myopic, F: full-knowledge)

Table 3: Percentage of optimal energy savings: varying α^d

Alg. \ α^d	0.0	22.2	44.4	66.7	88.9
M	82.2	82.4	83.0	84.0	84.0
P	85.1	85.4	85.7	86.2	86.4
F	86.4	86.4	86.4	86.4	86.4

(M: myopic, P: predictive non-myopic, F: full-knowledge)

can lead to significant energy reduction. This motivates considering more intelligent identification of key meetings to change their flexibility (described in the next section).

We also compared the performance of the three algorithms while varying the deadline flexibility, α^d . In Table 3, columns indicate different amounts of deadline flexibility and values are the optimality of each algorithm assuming a fixed time and location flexibility (α^T, α^L) = (67.5%, 47.1%). As we increase the deadline flexibility, both myopic and predictive non-myopic methods converge to the full-knowledge optimization result. This is because as the deadline flexibility increases, we can delay scheduling until we have more information. In our particular case, we do not necessarily see significant benefits by providing more deadline flexibility since the myopic and predictive non-myopic methods already achieved fairly high optimality compared to the full-knowledge method. While these percentages are small, given the vast amount of energy consumed by large-scale facilities, these reductions can lead to significant energy savings. We are investigating conditions where our algorithms get more benefits by deadline flexibility.

The same types of analysis are performed with another data set

Table 4: % of optimal energy savings: varying α^T and α^L

		Alg.		Location flexibility (α^L)		
				47.1	70.6	94.1
Time flexibility (α^T)	31.5	M	85.12	87.11	88.05	
		P	92.01	94.77	95.18	
		F	94.75	97.70	98.12	
	67.5	M	88.75	89.12	90.45	
		P	92.69	95.14	96.86	
		F	96.56	98.95	99.85	

(M: myopic, P: predictive non-myopic, F: full-knowledge)

Table 5: Energy improvement of identified key meetings (%)

$\alpha' \backslash \alpha$	(0,23.5)	(0,47.1)	(0,70.6)	(31.5,23.5)	(31.5,47.1)
(0,23.5)	-	-	-	-	-
(0,47.1)	15.78	-	-	-	-
(0,70.6)	30.00	27.89	-	-	-
(31.5,23.5)	30.75	-	-	-	-
(31.5,47.1)	44.25	35.25	-	28.87	-
(31.5,70.6)	45.25	37.12	32.16	30.15	25.78

from SMU and results were presented in Table 4. Please note that due to space limitation, we only show part of results in this paper.³ In the table, columns indicate the location flexibility (α^L) and rows indicate the time flexibility (α^T). We assume the deadline flexibility (α^d) of 0%. Similar to earlier results, the predictive method achieved about 97% optimality compared to the full-knowledge optimization and showed higher value than the myopic approach.

We investigated runtime comparisons of the three algorithms to verify the feasibility of our approaches to solve the real-world problem. The average runtime of the myopic optimization method was about 2 seconds, and the predictive non-myopic method was about 30 seconds.

4.1.3 Performance of identifying key meetings

We evaluated the performance of the algorithm to identify key meetings for energy reduction. In our tests, we selected 10 meetings and calculated the average energy savings if only one of these meetings changed their flexibility for each of the 10 selected meetings. This reflects an assumption that about 10% of people contacted will modify their flexibility.

Table 5 shows the average energy savings as described for various flexibility transitions. Columns indicate the initial level of flexibility ($\alpha = (\alpha^T, \alpha^L)$) and rows show the requested level of flexibility ($\alpha' = (\alpha'^T, \alpha'^L)$). For instance, the value in row 4 and column 3 (highlighted in the table) indicates a 27.89% average energy savings improvement if flexibility of 10 key meetings are changed from (0%, 47.1%) to (0%, 70.6%) is 27.89% one by one. An important interpretation of that results is that changing the flexibility of only one meeting (of 150 on average), when that one is from an appropriately chosen set contributed to significant energy savings. In the future, we will investigate another heuristic based on a learned profile of user likelihood of changing meeting flexibility.

We use the predictive non-myopic algorithm to identify key

³You can find the whole set of results here: <http://teamcore.usc.edu/junyounk/TESLA-sp.pdf>

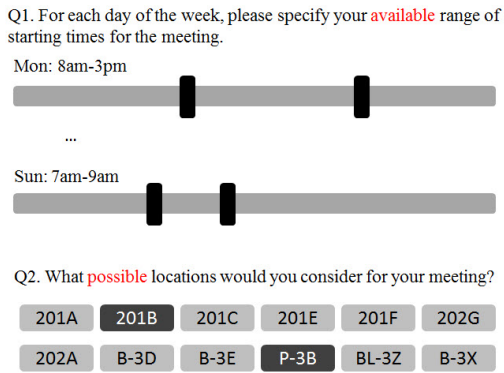


Figure 7: Screenshot of online survey




Figure 8: Diversity of people's flexibility

meetings, and we need to validate its accuracy. We checked the accuracy by directly comparing the meeting IDs of the key meeting set generated when using the predictive non-myopic and the full-knowledge optimization method. The average accuracy of our predictive method was over 93%, which supports that our detection algorithm is accurate.

4.2 Human Subject Experiments

The goal of human subject experiments is to support the results provided in the previous section by answering several questions: (i) are people flexible in real situations?; (ii) how flexible are they?; (iii) will people in the identified key meetings actually agree to change their flexibility to contribute energy savings?; and (iv) what would be a more effective way for an agent to persuade people? To answer these, we measure the amount of reported flexibility change while varying messaging.

We conducted two surveys on a pilot sample of participants (students on campus): (i) an online survey to understand flexibility of those who are using the testbed building; and (ii) a survey to measure flexibility change due to messaging.

4.2.1 Survey for initial flexibility

We conducted an online survey to understand the flexibility of meeting attendees (shown in Figure 7). The procedure to conduct this survey is as follows: we recruited 32 students who have used the meeting reservation system at the tested building and their facilities. We analyzed their profile including the details of their meeting requests and their flexibility in terms of time and locations considering their real constraints.

Figure 8 shows the distribution of the time and location flexibility. The x-axis shows the discretized flexibility level and their corresponding frequency in percentage is provided on the y-axis.

Table 6: Flexibility manipulation with various feedback (%)

	Group I	Group II
Average amount of change	5.15	17.12

People reported varied levels of time and location flexibility. The average time flexibility was 25.34% and their responses fell in a range of 9.86% and 42.86%. The average location flexibility was 16.05% and its range was 0 to 38.24%.

4.2.2 Survey for requested flexibility

We conducted a second survey to understand what types of feedback are most effective to change flexibility while scheduling meetings. We consider two test conditions: (i) feedback without motivation (Test Group I) (e.g., if necessary, do you think you will be able to provide more options in terms of time and location?), and (ii) feedback with motivation including average flexibility provided, and environmental motives (Test Group II) (e.g., on average, people who are using this system give 3–4 hour range for their available time on each day and 5–6 rooms for their available locations. This helps the system to compute more energy-efficient schedules that lead to energy savings by about 30% at the testbed building, which is equivalent to \$5,765 per year. Do you think you will be able to provide more options in terms of time and location?)

HYPOTHESIS 1. *More informed feedback (provided to subjects in Test Group II) will be more effective to conserve energy than feedback without motivation (Test Group I).*

To test the hypothesis above, we recruited 22 students with the same requirement of the earlier survey. Subjects were randomly tested under two different conditions when they accessed the online survey, and each test group had 11 individuals respectively.

Table 6 shows the average flexibility change in percentage (0–100%) of two test groups. When we provided more informed feedback including environmental motives (Group II), participants tripled their flexibility increase percentage. The difference is statistically significant and provides strong evidence for the hypothesis (t-test; $p < 0.01$).

This study shows that we can conserve energy by investigating methods to improve motivation to conserve energy by adjusting their flexibility.

5. RELATED WORK

TESLA is different from previous work by focusing on comfort-balanced energy-efficient incremental scheduling and identifying key meetings for adjusting their flexibility in commercial buildings. Furthermore, as an innovative application for energy savings, TESLA is evaluated on real meeting data (over 110,000 meetings and events) that have been collected from more than 500 rooms in nine educational buildings at USC and SMU. This combination of research contributions sets our work apart from previous research.

Energy Systems and Scheduling: Stein *et al.* [16] introduced a novel online mechanism that schedules the allocation of an expiring and continuously-produced resource to self-interested agents with private preferences while focusing on the fairness using pre-commitment in smart grid domain, which is not directly applicable in commercial buildings. There has been another work focusing on scheduling of home appliances considering user preferences [5, 15, 19], which is different from ours as preference considerations in their work is limited and we handle scheduling in large-scale commercial buildings.

Wainer *et al.* [17] presented a set of protocols for scheduling a meeting among agents that represent their respective user's interests and evaluated the suggested protocols while handling meeting scheduling problems. The objective in their work to find the optimal protocol to reach agreement among agents, which does not explicitly account for energy. Recently, there have been some other work considering meeting (re)location problems by exchanging messaging among agents [9, 10]. Although their work focused on minimizing energy consumption, they relied on the reactive scheduling and no flexibility model was considered.

There has been a significant amount of work done on online scheduling techniques to handle incremental requests considering temporal flexibility [8, 13]. Our work is different in focusing on energy-oriented scheduling in commercial buildings while allowing people to play a part in optimizing the operation in the building instead of managing the optimal resource allocation on buildings.

Social Influence in Human Subject Studies: Wood and Neal [18] have studied the potential of interventions to reduce energy consumption and they have shown that it is not only to change workplace energy consumption but also to establish energy use habits that maintain over time.

In social psychology, there has been a significant deal of work to figure out the correlation between irritation/distraction factors and persuasion. McCullough and Ostrom [11] and Cacioppo and Petty [7] discussed that message repetition would increase positive attitudes in a situation where highly similar communications are used and showed that there is a positive relationship between the number of presentations and attitude from general social psychology perspectives. Focusing on a commercial advertisement, Pechmann and Stewart [12] predicted the effectiveness of different strategies on advertising and examined the effects of message repetition on attitude changes. In addition, Baron *et al.* [6] discussed that distractions affect behavior decisions, but they are more or less effective in increasing persuasion depending upon whether people can easily ignore the distraction.

We leverage insights from social psychology in understanding and designing reliable and accurate human behavior models.

6. CONCLUSION

The key contribution of this innovative application paper is not just our agent TESLA, but more importantly, TESLA's analysis of real-world data — 32,000 meetings from the University of Southern California (USC), and 80,000 meetings from the Singapore Management University (SMU) — to show the power of flexibility. TESLA's promise of energy savings is rooted in this real data, and illustration that significant energy savings may accrue not from imposing any complex interaction protocol on humans, but from a simple action of providing schedule flexibility. More specifically, TESLA provided three key contributions. First, it provided three online scheduling algorithms that consider the diversity of people's flexibility for energy-efficient scheduling of incrementally/dynamically arriving meetings and events. Second, it presented an algorithm to effectively identify key meetings that lead to significant energy savings by adjusting their flexibility. Lastly, surveys of real users were provided indicating that TESLA's savings can be realized in practice. We showed that, compared to the current systems, TESLA can substantially reduce the overall energy consumption.

7. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1231001. Any opinions, find-

ings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

8. REFERENCES

- [1] Average energy prices in the los angeles area. http://www.bls.gov/ro9/cpilosa_energy.htm.
- [2] Efficient class scheduling conserves energy. <http://goo.gl/cZwgB>.
- [3] New classroom scheduling methods save energy, money for msu. <http://news.msu.edu/story/6501/>.
- [4] Meetings in america: A study of trends, costs, and attitudes toward business travel and teleconferencing, and their impact on productivity. 2001.
- [5] T. Bapat, N. Sengupta, S. K. Ghai, V. Arya, Y. B. Shrinivasan, and D. Seetharam. User-sensitive scheduling of home appliances. In *SIGCOMM*, 2011.
- [6] R. Baron, P. Baron, and N. Miller. The relation between distraction and persuasion. *Psychological Bulletin*, 80(4):310, 1973.
- [7] J. Cacioppo and R. Petty. Effects of message repetition on argument processing, recall, and persuasion. *Basic and Applied Social Psychology*, 10(1):3–12, 1989.
- [8] A. Gallagher, T. Zimmerman, and S. Smith. Incremental scheduling to maximize quality in a dynamic environment. In *ICAPS*, 2006.
- [9] J. Kwak, P. Varakantham, R. Maheswaran, M. Tambe, T. Hayes, W. Wood, and B. Becerik-Gerber. Towards robust multi-objective optimization under model uncertainty for energy conservation. In *AAMAS Workshop on Agent Technologies for Energy Systems (ATES)*, 2012.
- [10] J. Kwak, P. Varakantham, R. Maheswaran, M. Tambe, F. Jazizadeh, G. Kavulya, L. Klein, B. Becerik-Gerber, T. Hayes, and W. Wood. SAVES: A sustainable multiagent application to conserve building energy considering occupants. In *AAMAS*, 2012.
- [11] J. McCullough and T. Ostrom. Repetition of highly similar messages and attitude change. *Journal of Applied Psychology*, 59(3):395, 1974.
- [12] C. Pechmann and D. Stewart. Advertising repetition: A critical review of wearin and wearout. *Current issues and research in advertising*, 1988.
- [13] N. Policella, S. F. Smith, A. Cesta, and A. Oddi. Incremental scheduling to maximize quality in a dynamic environment. In *ICAPS*, 2004.
- [14] P. Scerri, D. V. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *JAIR*, 17:171–228, 2002.
- [15] K. C. Sou, J. Weimer, H. Sandberg, and K. H. Johansson. Scheduling smart home appliances using mixed integer linear programming. In *CDC-ECC*, 2011.
- [16] S. Stein, E. Gerding, V. Robu, and N. Jennings. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *AAMAS*, 2012.
- [17] J. Wainer, P. R. F. Jr., and E. R. Constantino. Scheduling meetings through multi-agent negotiations. *Decision Support Systems*, 44(1), 2007.
- [18] W. Wood and D. Neal. The habitual consumer. *Journal of Consumer Psychology*, 19:579–592, 2009.
- [19] G. Xiong, C. Chen, S. Kishore, and A. Yener. Smart (in-home) power scheduling for demand response on the smart grid. In *ISGT*, 2011.